


# ZUPS Sequence Timing

Wed, Dec 31, 2003



A series of 8 consecutive 15 Hz cycles is shown. In the first sequence, the normal STT query is sent to each power supply to get the latest readings, settings, and status information. It normally takes about 114 ms from the time the STT command is sent until the reply is received. For this reason, successive power supplies are addressed every 2 cycles.

The second sequence illustrates the situation when a reply is missing. After the timeout shown, the address is sent for the next supply, and on the following cycle the STT query is sent to that supply. Because of the timeout, 2 cycles are lost awaiting the reply in vain.

The third sequence shows what happens following a timeout. On the next cycle in which it is the failed power supply's turn, the REV command is sent to get the peak voltage and current and the software version number. (Since that reply is only 14 characters, it is returned much more quickly than the 57-character STT reply.) When the version number has been received and interpreted, the normal STT command is sent to the same supply, now that it has shown itself to be "up" again.

Each command is sent early in the 15 Hz cycle, about 1–2 ms after the start of the cycle. The query commands consist of 6 characters, such as ":STT?;". The STT reply consists of 57 characters. At 9600 baud, or about 1 ms per character, the serial bandwidth time is therefore about 63 ms. Since the total elapsed time is 114 ms, about 50 ms is needed for the supply to construct its reply. As another example, the REV command returns only 14 characters in 24 ms, so the power supply needs  $24 - (6 + 14) = 4$  ms to construct that reply. The RMT command returns a short reply of only 4 characters after only 11 ms, implying 1 ms for processing. Upon successful processing of the STT command reply at about 50 ms past the start of the cycle, the 7-character ADR command is sent to select the next power supply in sequence. The command sent to that supply is sent at the start of the following cycle.

Each communication with a supply ends with the STT query. It is only preceded by another command if the time is "ripe" for one. When a supply is considered "down" because it failed to reply to the previous query, the REV command is sent in place of the STT command. Only if the REV command reply is forthcoming does it get the STT command. In this way, the sequencing proceeds through all supplies, only pausing to send one extra command, if needed. With 4 power supplies, the usual sequencing of STT commands takes 8 cycles, as illustrated above.

After some study of unexpected failures to receive replies, the program logic was changed to allow an extra cycle of delay following the sending of the ADR command to select a new supply. After this change, the errors ceased, implying that extra time is needed after an ADR before sending another command. The normal sequence timing now looks like the first example above, except that an extra cycle of delay occurs after selecting the next power supply. As a result, it now takes at least 3 cycles to collect the latest data from each supply. The complete "tour" through 4 supplies therefore takes 0.8 seconds.